UC Merced
Proceedings of the Annual Meeting of the Cognitive Science Society

Title
The role of contrast in category learning

Permalink
https://escholarship.org/uc/item/9511v86v

Journal
Proceedings of the Annual Meeting of the Cognitive Science Society, 47(0)

Authors
Altenhof, Aja Marie
Lupyan, Gary

Publication Date
2025

Copyright Information
This work is made available under the terms of a Creative Commons Attribution License, available at https://creativecommons.org/licenses/by/4.0/

Peer reviewed

eScholarship.org Powered by the California Digital Library
University of California

The role of contrast in category learning

Aja Altenhof (altenhof@wisc.edu)
Department of Linguistics
University of Wisconsin, Madison

Gary Lupyan (lupyan@wisc.edu)
Department of Psychology
University of Wisconsin, Madison

Abstract

Word meanings are contrastive. When we are told that something is a square, we are also told that it is not a triangle. However, words may be learned in different contrasts. One person might learn about squares in contrast with circles. Does this mean that the two have different representations of "square?" To answer this question, participants learned to label novel shapes with novel labels in a category learning task. Critically, we manipulated the contrast participants received during learning: an A-shape is specifically not a B or an A-shape is specifically not a D. Afterwards, we tested participants' knowledge of the learned categories using explicit categorization tasks and similarity judgments. Contrast during learning mattered. Shapes from contrasted categories were categorized more accurately, were less confusable and rated as less similar.

Keywords: categorization; contrast; word learning; similarity

Introduction

Many word meanings are partly defined by what they aren't. A weekend is not a weekday. A suburb is an area that is neither urban nor rural. Vegetarians are people who don't eat meat. In these cases, contrast doesn't merely complement meaning—it actively constructs it. /

Language users are attuned to this idea of definition-by-contrast. For example, when asked to generate semantic associates for words, participants often produce members of a contrastive category, like "cat" for "dog" (G. L. Murphy & Andrew, 1993). Likewise, when asked to list items that one usually finds in a kitchen, people–not surprisingly–list typical items like a stove and a refrigerator. But when asked to list what is *not* found in a kitchen, people also produce similar responses, listing items like "toilet" and "bed"–items typical of categories that are in direct contrast with "kitchen" (Greene, 2016). A kitchen isn't solely defined by what it contains, but in opposition to similar spaces.

In structuralist theories of semantics, words don't exist in isolation. Instead, meaning is derived through the relationship to other words and concepts in a network (De Saussure, 1916). Under this framework, words are organized into interconnected fields or domains, where meaning emerges through opposition (M. L. Murphy, 2003).

Given that words can be defined in relation to what already exists, learners expect new words to be different. A new concept could be representative (Tenenbaum & Griffiths, 2001) of the opposite category (e.g., the best dog is the least cat-like dog) or maximally dissimilar (Austerweil, Liew, Conaway, & Kurtz, 2024) from an existing one (e.g., the best dog is the most dog-like dog). This tendency extends to how learners represent and understand categories: when participants learn categories that contrast on specific dimensions, their representations become idealized, exaggerating the differences along contrasted dimensions (Davis & Love, 2010). Contrast doesn't just refine existing categories but fundamentally guides the creation of new concepts.

Direct contrast between concepts may be more effective than learning them in isolation. Building on the distinction between blocked and interleaved learning, Kattner, Cox, and Green (2016) systematically compared multiple learning approaches to understand how each approach affects category representations. Participants were trained on four novel categories using one of four tasks: blocked learning, interleaved learning, identification training (answering yes/no if a stimulus matched a category label), and triplet training (matching one of three labels to a stimulus). Interestingly, the most effective approach directly leveraged contrast. Of the four tasks, identification training, which explicitly required participants to judge what categories didn't include, transferred most successfully to a novel discrimination task. By drawing attention to the relationships between categories, identification training led to more robust outcomes.

Definition-by-contrast is further complicated by the internal structure of categories themselves. Categories are not uniform: some members are more typical or representative than others (Rosch & Mervis, 1975). These prototypical members share more features with other category members and fewer features with members of contrasting categories. As a result of this graded structure, more central members show robust processing advantages: they are categorized more quickly (Thorpe, Fize, & Marlot, 1996), recognized more accurately (G. L. Murphy & Brownell, 1985) and serve as cognitive reference point for the category as a whole (Rosch & Mervis, 1975).

The competition between typicality and contrast creates an interesting explanatory tension. While typicality effects suggest that category representations are organized around a central conceptual core, contrast effects suggest that categories are defined by boundaries. Given this asymmetry, contrast could disproportionately affect boundary members, sharpening distinctions where categories are most confusable; learning that 'A is not B,' might be most useful for atypical A

In D. Barner, N.R. Bramley, A. Ruggeri and C.M. Walker (Eds.), *Proceedings of the 47th Annual Conference of the Cognitive Science Society* ©2025 the author(s). This work is licensed under a Creative Commons Attribution 4.0 International License (CC BY).

exemplars that share features with B. Alternatively, contrast could accentuate diagnostic features, strengthening the distinctiveness of typical members; learning 'A is not B' could reinforce the characteristics of both categories.

The function of contrast may depend on category structure, the relevant dimension, and task demands. Perceptual discrimination between similar stimuli increases along categorization-relevant dimensions and is particularly pronounced at category boundaries(Goldstone, 1994). But while some dimensions primarily show assimilation, or increased similarity effects, others primarily show contrast effects, depending on the extremeness of the stimuli being compared (Barker & Imhoff, 2021). /

The current experiment explores how contrast shapes category representations. Participants were taught four novel shape categories paired with four novel labels (for clarity, we will refer to them 'A', 'B', 'C' and 'D') arranged in a circular stimuli space that varies on a single continuous dimension. Though adjacent categories in this space are theoretically equally perceptually similar (As are as similar to Bs as they are to Ds), participants were assigned to different learning conditions, manipulating which category was contrasted with the target. For example, on a training trial, a person assigned to the AB/CD contrast saw the A label paired with an A shape (the correct response) and a B shape (the incorrect response). A person assigned to the AD/BC contrast would see an A label paired with the same A shape, but would see a D shape as the incorrect option. So for AB/CD participants, an A is not just an A, but also explicitly not a B. For AD/BC participants, an A is explicitly not a D.

After learning, participants completed three tasks to assess their acquisition and representation of the new categories: matching the shape to its name (Label-Choice), similarity judgments, and matching a name to a shape (Shape-Choice). Drawing on findings that contrast supports category distinction and generalization, we expect that category pairs will be perceived differently based on their training history. Specifically, categories that were directly contrasted during learning (e.g., learning that A is not B) should be rated as less similar than pairs that were never contrasted (e.g., never learning about A and C together). Furthermore, we expect contrast to interact with typicality. If contrast sharpens category boundaries, we should observe a stronger benefit from contrast for boundary exemplars than for typical ones—potentially closing the typicality gap in categorization performance (Rosch & Mervis, 1975). Alternatively, if contrast augments within-category similarities, we could see enhanced typicality effects, where typical members are even more easily distinguished from those in contrasted categories.

Methods

Participants

We recruited 223 participants through the UW-Madison SONA participant pool, who received course credit for their time.

Stimuli

We used the "Validated Circular Shape" (VCS) space developed by Li, Liang, Lee, and Barense (2020). The VCS is intended to mirror the color wheel, such that angular distance along a 2D circle reflects visual similarity. This creates a continuous space of ostensibly perceptually equidistant shapes. As a result, a given shape is equally similar to the shape the same number of degrees away in either direction on the circle. Although the shape space was previously validated through human similarity ratings (Li et al., 2020), as we shall see, there is reason to think the space is not in fact perceptually balanced, which required us to conduct an additional norming study.

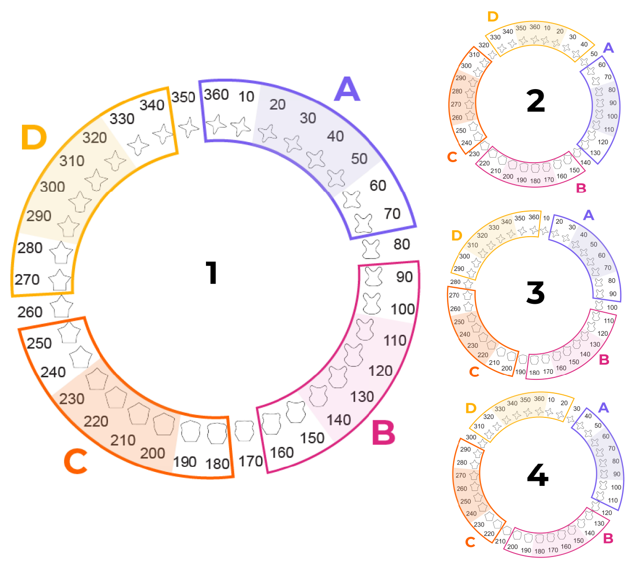

Figure 1: The category space in all conditions (1, 2, 3 and 4). Given the target A category, if B serves as the "contrast," then D is the "neighbor" and C is the "orthogonal" category.

We divided the VCS space into four categories of eight shapes each (Figure 1). Within each category, shapes were divided into inner members (highlighted), which are more central and so more typical members of the category, and outer members. The outer members were located closer to category boundaries and thus were more similar to the adjacent categories than the inner shapes. The use of normed stimuli aimed to ensure that differences in participants' category representations were not due to differences in perceptual distance between categories, however Zettersten, Suffill, and Lupyan (2020) found that the visual discriminability of these shapes was moderated by nameability. To mitigate this and ensure generalizability, participants were randomly assigned to one of four conditions that corresponded to different category configurations in the space. Conditions were created by shifting the categories 20 degrees clockwise. Each category was given one of four novel labels: *toma*, *ridu*, *loti* or *fimo*. Label assignment was counterbalanced across participants.

Procedure

The experiment was divided into four tasks (Figure 2): initial learning, label-choice, similarity-judgments, and generalization.

Learning During learning, participants were shown two shapes from different categories and a category label. They were instructed to select the correct shape for the label using the 'z' or '/' keys, for the left and right shape respectively. After each trial, participants received auditory feedback.

Though participants were trained on all four categories, they were randomly assigned to one of two conditions manipulating the pairs of categories shown together, or contrasted, on a given learning trial (Figure **??**): AB/CD (A is contrasted with B and C is contrasted with D) or AD/BC (A is contrasted with D and B is contrasted with C). The specific contrast established during learning creates distinct relationships between the categories, allowing us to examine how contrast affects accuracy, generalization and perceived similarity. For a participant in the AB/CD condition, if Category A was the target shape, Category B would be the contrasted category (adjacent to A and directly paired with A during learning). Category D would be the neighboring category (adjacent to A but never directly contrasted with it), and Category C would be the orthogonal category (opposite to A).

Crucially, only inner category members (i.e., the more typical shapes) were presented in learning trials. This enabled us to test how well participants generalized to the outer (less typical) shapes, and how typicality might interact with contrast.

Label-Choice After learning, participants proceeded to the Label-Choice task where they were asked to match a shape to the appropriate label. Participants did not receive feedback during this task. Participants completed 32 trials: one for each shape.

Similarity Judgment Following the Label-Choice task, participants completed similarity judgments. They were instructed to select which of three shapes in a row was most similar to a target shape displayed above them. Each row contained one shape from the contrasted, orthogonal, and neighboring categories, relative to the target shape, with their positions randomized across trials. All similarity judgments used only inner (typical) category members.

To control for the effect of perceptual distance, the contrast and neighbor shapes were the same number of degrees away from the target. For example, in the AB/CD contrast condition, if the target was from Category A at 50 degrees, both the neighboring category shape (from Category D) and the contrasted category shape (from Category B) would be positioned 90 degrees away from the target. The orthogonal category shape was always positioned 180 degrees away from the target.

Shape-Choice The structure of the last task was the same as learning. However, this task differed in two critical ways: participants received no feedback, and shape pairs were no longer restricted to contrasted categories or typical exemplars. Instead, trials included all possible shapes (both inner/typical and outer/atypical members) and all possible category pairings (contrasted, neighboring, and orthogonal combinations). Including combinations not seen during learning allowed us to assess how well participants generalize their experience learning contrasted categories to the rest of the space.

Results

Data are available here. Analyses were performed using R (R Core Team, 2013) and graphs were created using ggplot2 (Wickham, 2011).

Learning Overall accuracy was high ($M = 0.90, SD = 0.29$), reaching 96 % ($SD = 0.19$) by the final 25% of trials, indicating participants successfully learned the categories. 3 participants were excluded for missing data. 18 participants with mean accuracies below 0.75 were excluded for failure to learn.

Label-Choice In the Label-Choice task, participants needed to match a shape to the correct label. Participants were mostly accurate ($M = 0.76, SD = 0.43$), even for outer shapes ($M = 0.69, SD = 0.46$), which they hadn't seen during training. Accuracy during this task was correlated with participants' learning; better learners were more accurate overall, $\rho(201) = 0.30, p < .001$, and better able to generalize to less typical category members, $\rho(201) = 0.27, p < .001$.

More interesting is what happened when participants made errors (Figure 3). We conducted a linear regression to understand which incorrect category label (contrast, neighbor, or orthogonal) participants selected when they didn't choose the target one. Unsurprisingly, orthogonal category labels, representing shapes furthest from the target in the stimuli space, were selected significantly less frequently than the contrastive category baseline ($b = -0.30, SE = 0.03, t(585) = -11.81, p < .001$). More importantly, neighboring category labels were selected significantly more frequently than contrastive ones ($b = 0.15, SE = 0.03, t(585) = 5.91, p < .001$), despite both categories being the same angular distance from the target shape. Consistent with our predictions, contrasted categories became less confusable as a result of their explicit pairing during learning.

Similarity Judgments In the Similarity Judgment task, participants indicated which of three shapes (representing the contrast, neighbor, and orthogonal categories) was most similar to the top shape. As expected, a linear regression revealed that the orthogonal category was chosen as "most similar" significantly less frequently than either the neighbor or contrast categories ($b = -0.44, SE = 0.03, t(600) = -14.18, p < .001$), which were selected at similar rates ($b = 0.04, SE = 0.03, t(600) = 1.32, p = .188$).

However, closer analysis revealed a strongly bimodal distribution (Figure 4): approximately half of the participants

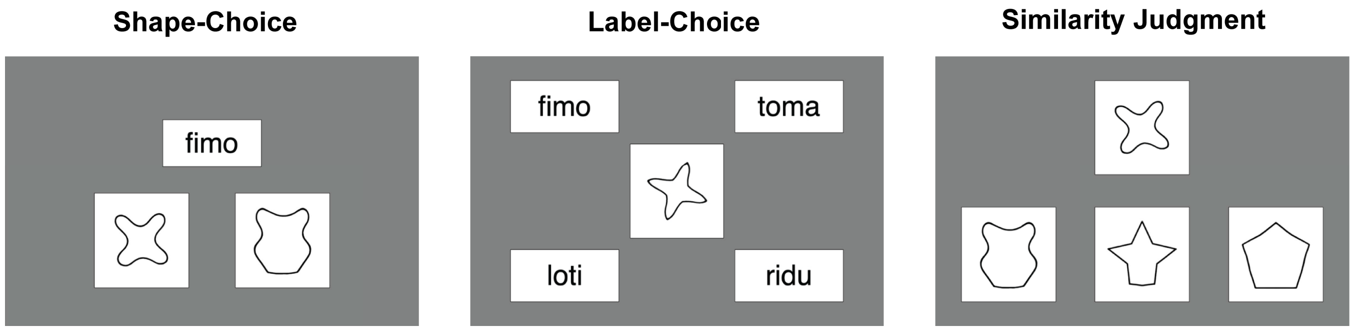

Figure 2: The four experimental tasks: Shape-Choice (Learning and Generalization), Label-Choice, and Similarity Judgments. In the Shape-Choice tasks, participants needed to match the correct shape to the label. In the Label-Choice task, participants needed to pick the right label for the given shape. Finally, in the Similarity Judgment task, participants selected which shape was most similar to the top one.

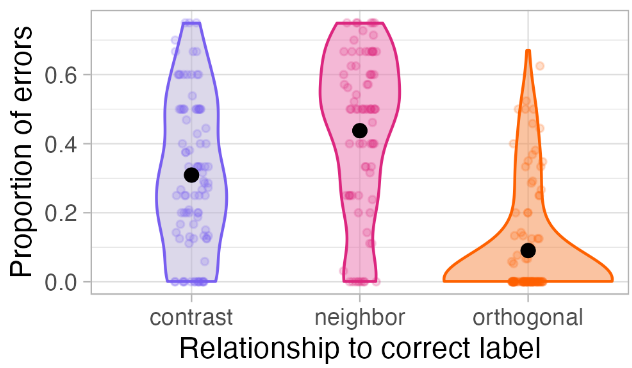

Figure 3: Distribution of participant errors in the Label-Choice task by category relationship

consistently chose the contrastive category ($M = 0.46, SD = 0.38$), while the others consistently selected the neighbor ($M = 0.50, SD = 0.38$). Further inspection showed that these preferences were strongly predicted by participants' assigned stimulus condition (1, 2, 3 or 4) and contrast (AB/CD or AD/BC), suggesting the shape space was not perceptually well-equated as intended. To address this, we collected baseline similarity data from a separate group of 91 MTurk participants who completed the similarity judgment task without any category training. We then used these baseline similarities as a covariate to control for the inherent perceptual biases when analyzing the likelihood of contrast versus neighbor choices.

For the analyses that follow, we exclude the predictably rare "orthogonal" responses and include just the trials on which participants chose the "neighbor" or "contrast," to better understand what guided these choices. We ran a linear mixed-effects model predicting participants' shape selection with fixed effects for baseline probability, contrast type (AB/CD vs. AD/BC), nameability, stimulus type (neighbor vs. contrast), and their interactions. We also included random slopes for stimulus type by participant.

The direction of contrast (AB/CD or AD/BC) had the largest effect on participants' similarity judgments. For the AD/BC training condition, participants were significantly more likely to choose neighbor stimuli, while for the AB/CD training condition, participants showed a strong preference for contrast stimuli ($\beta = 0.29, t = 26.06$). This suggests that, regardless of contrast, participants had a preference for grouping portions of the circle together, such that right/left groupings were preferable to top/bottom ones. That is, A's were always chosen as more similar to Bs than Ds (and C's to Ds than Bs).

We also found that baseline similarity was a predictor of which supposedly perceptually equidistant shape participants chose ($\beta = -0.018, t = -14.68$). If indeed angular distance corresponded to perceptual distance as you traverse the space in either direction, there should have been no systematic baseline preference between shapes equidistant from the target. This finding suggests that the VCS stimuli are not as well-normed as described in Li et al. (2020).

Contrast and baseline similarity further interacted with whether the shape was the contrast or neighbor ($\beta = -0.011, t = -8.93$). Though baseline similarity had a stronger influence on AB/CD contrast participants' choices overall, this effect was more pronounced for contrast shapes than for neighbor ones. Conversely, for participants who received an AD/BC contrast, baseline similarity was weaker in general and most pronounced for neighbor stimuli, suggesting that the impact of contrast depends on prior similarity. More specifically, when combined with participants' general lateral similarity preferences, learned contrast was most effective when it supported an existing perceptual contrast, dividing the space horizontally.

Shape-Choice This task mirrored learning—participants needed to match the name to the correct shape—but included all contrasts (AB, BC, AD, AC, etc.) and all shapes, both inner (more typical) and outer (less typical).

Mirroring the Label-Choice task, overall accuracy was

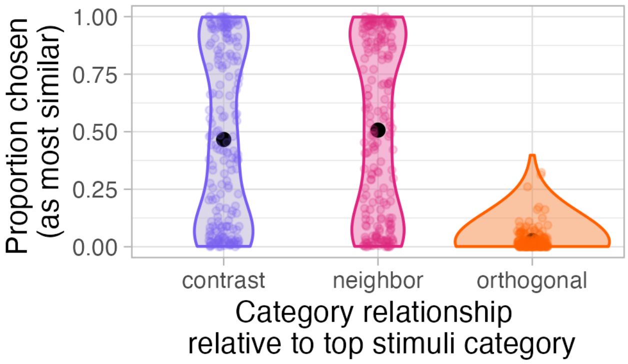

Figure 4: Distribution of participant choices as "most similar" in the Similarity Judgment task by category relationship

high ($M = 0.90, SD = 0.30$) and correlated with learning, $\rho(201) = 0.47, p < .001$. To examine the effect of category relationship on accuracy, we fit a generalized linear mixed-effects model with relationship type as a fixed effect and random intercepts for participants. Unsurprisingly, participants were best at trials between shapes from orthogonal categories ($M = 0.92, SD = 0.26$), which by nature compared shapes further away in the category space than neighbor and contrast trials ($\beta = 0.24, z = 6.42, p < .001$).

Focusing on just adjacent (contrast and neighbor) trials, there were significant main effects of relationship type, stimuli location, and angular distance (Figure 5). Participants were less accurate on neighbor trials compared to contrast trials ($\beta = -0.18, z = -4.70, p < .001$), suggesting that training enhanced discrimination for directly contrasted categories. Similarly, participants were less accurate at trials that contained less typical outer shapes ($\beta = -0.19, z = -7.94, p < .001$), replicating classic typicality findings. We also found an interaction between trial type and typicality ($\beta = 0.04, z = 2.61, p = .009$): while participants generally performed worse on neighbor trials and with atypical shapes, the negative effect of neighbor trials was reduced when involving atypical shapes. In other words, the accuracy advantage for contrast trials (which participants had been trained on) was more pronounced with typical shapes and diminished with atypical ones. The benefits of contrastive training may be reduced when learners are asked to generalize to unseen exemplars.

Discussion

How much of knowing something is also knowing what it's not? In this experiment, we investigated how explicit contrast provided during learning influences category representations. Participants were trained on four novel shape categories, learned in contrast with each other (e.g., "an A is not a B" or "an A is not a D"). Afterwards, we evaluated participants' category representations using both categorization accuracy (in the Shape-choice and Label-Choice tasks) and explicit similarity judgments to examine how contrast shapes

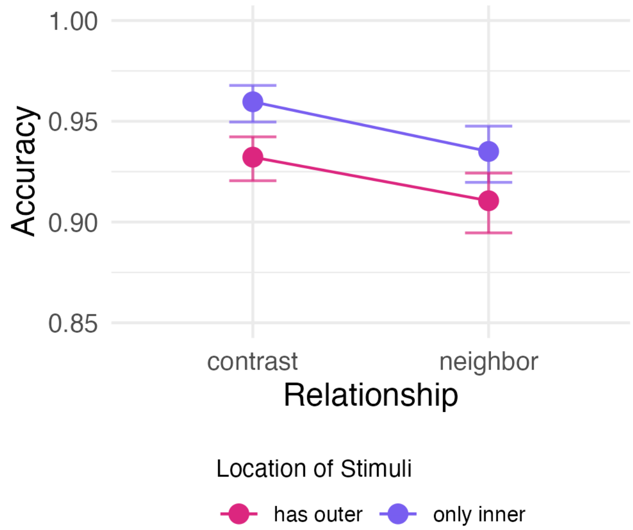

Figure 5: Participant accuracy in Shape-Choice Task by category relationship and stimuli typicality

category boundaries, affects typicality gradients, and influences generalization to novel examples.

As predicted, contrast during learning mattered. Across two tasks, participants showed better discrimination between categories that had been directly contrasted during training compared to those that were never contrasted. In the Label-Choice task, participants were less likely to confuse the label for a target shape with the label for the contrastive category, suggesting that contrast enhanced between-category distinctiveness. This result is particularly striking given that the labels for the contrastive shapes consistently co-occurred during learning; the label for A was always seen with B shapes, and vice versa. While previous work (Zettersten et al., 2020) has demonstrated that object co-occurrence and overlapping visual contexts, increased confusability—object associations came at the cost of successful word learning—we find the opposite. Rather than making categories more confusable through association, contrastive co-occurrence appeared to facilitate their discrimination.

The Shape-Choice task showed a similar pattern of results; participants were more accurate on contrast-typical trials than neighbor-typical trials, despite both trial types displaying stimuli that appeared equally often during training. This suggests an advantage when matching labels to referents in familiar contrastive contexts, reinforcing the finding that contrast during learning improves discriminability.

Beyond the benefit of contrast, we also observed a typicality advantage: participants were more accurate at categorizing inner (typical) shapes on both explicit categorization tasks, consistent with classic literature (Jolicoeur, Gluck, & Kosslyn, 1984). However, the interaction between contrast and typicality in the Shape-Choice task revealed a more nuanced picture. If contrast sharpens category boundaries, we would expect greater benefits for atypical or ambiguous shapes near these boundaries, where discrimination is most challenging. Instead, we found the reverse: contrast supported the categorization of typical shapes. Though contrast didn't preferentially reinforce category boundaries, it's unclear if this benefit for typical shapes was driven by the alternative—emphasizing category similarities—or straightforward familiarity. Participants may have performed better simply because they saw only typical examples in contrastive frameworks during training.

This confound between a focus on shared features and familiarity leaves the mechanism that drives contrast an open question. Two alternative (but potentially simultaneous) accounts could explain how contrast benefits categorization. First, by accentuating the differences between in-group and out-group members, contrast might push categories apart on the most relevant dimension (Davis & Love, 2010), effectively sharpening category boundaries similar to Tenenbaum and Griffiths' (2001) idea of representativeness. Second, by drawing attention to shared features within a category, contrast could create tighter, more cohesive categories through compression. Individual exemplars could become more similar to each other, highlighting central category features in a manner similar to Austerweil et al.'s (2024) maximal dissimilarity principle. While our results tentatively challenge the boundary-sharpening (distance-based) account, the current experiment was not specifically designed to distinguish between these hypotheses.

In addition to improving categorization and generalization, contrast during learning also impacted similarity judgments; participants who saw two categories contrasted during learning rated them as less similar, compared to baseline participants who received no category training. Even though contrasted categories always occurred in the same contexts, they were chosen as less similar to a category that was never seen with the target, the neighbor.

The direction of contrast (AB/CD or AD/BC) was the primary driver of participants' choices. Learners who received an AB/CD contrast consistently rated shapes from the contrasted category as more similar, while those who received an AD/BC contrast preferred neighbor shapes. This further interacted with baseline similarity, which had a stronger influence on participants who received an AB/CD contrast compared to those in the AD/BC condition. That is, contrastive learning more strongly overrode learners' pre-existing perceptual biases for AD/BC participants. The effect of contrast may depend significantly on prior similarity.

Why are some contrasts more effective than others? Definition by contrast is not only about dissimilarity. While a concept is technically not many things (a dog is not a cat, but it's also not a book or a mountain), only some of these contrasts are relevant. The most functional contrasts are between concepts that frequently co-occur or serve similar functions. The things that something "is not" may still be quite similar. When asked to list what is not found in a kitchen, participants don't mention logically correct but contextually irrelevant items like airplanes or volcanoes. Instead, they generated items like beds and toilets—objects that serve similar domestic functions, and belong to the same superordinate category, but are different on a specific dimension (Greene, 2016).

It's possible that contrast works best when emphasizing existing distinctions, instead of creating them from scratch. Though shapes in category A were supposed to be equally similar to both Bs and Ds, our results suggest this wasn't actually the case. The asymmetry in participants' responses between the AB/CD and AD/BC contrasts, plus the influence of baseline similarity ratings, indicates that the VCS space had inherent perceptual groupings. Specifically, categories seemed more similar when grouped laterally rather than horizontally. This means some category boundaries (like between A and D) might have been more naturally salient than others (like between A and B). If indeed A's are more similar to Bs than Ds, participants could've leverage the existing dissimilarity between As and Ds during contrastive learning. Contrast, augmented with a linguistic cue, may have reinforced an existing perceptual boundary.

Though the VCS space did not function as intended, it nonetheless revealed important insights about the role of contrast. Ultimately, in real-world category learning, objects aren't perceptually equidistant on a single continuous dimension. Natural categories have complex and overlapping structures that vary on multiple features—no single feature makes a cat not a dog. Our finding that contrast interacts with similarity may more accurately reflect actual categorization.

Future work can better account for this relationship between contrast and pre-existing perceptual organization. Controlling for underlying forms of similarity will allow us to better understand what contrast is doing, whether it's sharpening boundaries, tightening categories, or creating new distinctions between them.

Our results reveal a complex relationship between contrast and category learning. Though contrast made categories less similar and more distinguishable in explicit categorization tasks, what something is contrasted with—and its relative similarity—seems to matter.

References

Austerweil, J. L., Liew, S. X., Conaway, N., & Kurtz, K. J. (2024, July). Creating Something Different: Similarity, Contrast, and Representativeness in Categorization. *Computational Brain & Behavior*. doi: 10.1007/s42113-024-00209-5

Barker, P., & Imhoff, R. (2021, November). The dynamic interactive pattern of assimilation and contrast: Accounting for standard extremity in comparative evaluations. *Journal of Experimental Social Psychology*, *97*, 104190. doi: 10.1016/j.jesp.2021.104190

Davis, T., & Love, B. C. (2010). Memory for category information is idealized through contrast with competing options. *Psychological Science*, *21*(2), 234–242.

De Saussure, F. (1916). Nature of the linguistic sign. *Course in general linguistics*, *1*, 65–70.

Goldstone, R. L. (1994). Influences of categorization on perceptual discrimination. *Journal of Experimental Psychology: General*, *123*(2), 178–200. doi: 10.1037/0096-3445.123.2.178

Greene, M. R. (2016). Estimations of object frequency are frequently overestimated. *Cognition*, *149*, 6–10.

Jolicoeur, P., Gluck, M. A., & Kosslyn, S. M. (1984). Pictures and names: Making the connection. *Cognitive psychology*, *16*(2), 243–275.

Kattner, F., Cox, C. R., & Green, C. S. (2016, October). Transfer in Rule-Based Category Learning Depends on the Training Task. *PLOS ONE*, *11*(10), e0165260. doi: 10.1371/journal.pone.0165260

Li, A. Y., Liang, J. C., Lee, A. C. H., & Barense, M. D. (2020, May). The validated circular shape space: Quantifying the visual similarity of shape. *Journal of Experimental Psychology: General*, *149*(5), 949–966. doi: 10.1037/xge0000693

Murphy, G. L., & Andrew, J. M. (1993). The conceptual basis of antonymy and synonymy in adjectives. *Journal of memory and language*, *32*(3), 301–319.

Murphy, G. L., & Brownell, H. H. (1985, January). Category differentiation in object recognition: typicality constraints on the basic category advantage. *Journal of Experimental Psychology. Learning, Memory, and Cognition*, *11*(1), 70–84. doi: 10.1037//0278-7393.11.1.70

Murphy, M. L. (2003). *Semantic relations and the lexicon: Antonymy, synonymy and other paradigms*. Cambridge University Press. (Google-Books-ID: 7pAIpz87jbEC)

Rosch, E., & Mervis, C. B. (1975). Family resemblances: Studies in the internal structure of categories. *Cognitive Psychology*, *7*(4), 573–605. doi: https://doi.org/10.1016/0010-0285(75)90024-9

Team, R. C. (2013). *R: A language and environment for statistical computing*. Vienna, Austria: R Foundation for Statistical Computing. Retrieved from `http://www.R-project.org`

Tenenbaum, J. B., & Griffiths, T. L. (2001, August). Generalization, similarity, and Bayesian inference. *Behavioral and Brain Sciences*, *24*(4), 629–640. doi: 10.1017/S0140525X01000061

Thorpe, S., Fize, D., & Marlot, C. (1996). Speed of processing in the human visual system. *nature*, *381*(6582), 520–522.

Wickham, H. (2011). ggplot2. *Wiley interdisciplinary reviews: computational statistics*, *3*(2), 180–185.

Zettersten, M., Suffill, E., & Lupyan, G. (2020). Nameability predicts subjective and objective measures of visual similarity. *Proceedings of the 42nd Annual Conference of the Cognitive Science Society*.